

UNITED STATES PATENT APPLICATION

For

METHOD AND APPARATUS FOR MANAGING A
POWER LOAD CHANGE IN A SYSTEM

INVENTOR:

William P. Cornelius

Prepared by:

BLAKELY SOKOLOFF TAYLOR & ZAFMAN LLP
12400 Wilshire Boulevard
Los Angeles, CA 90025-1026
(408) 720-8300

"Express Mail" mailing label number: EL 672 749 483 US

Date of Deposit: January 18, 2002

I hereby certify that I am causing this paper or fee to be deposited with the United States
Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that
this paper or fee has been addressed to the Commissioner for Patents, Washington, D. C. 20231

Connie Thayer

(Typed or printed name of person mailing paper or fee)

Connie Thayer

(Signature of person mailing paper or fee)

1-18-2002

(Date signed)

04860.P2676

METHOD AND APPARATUS FOR MANAGING A POWER LOAD CHANGE IN A SYSTEM

FIELD OF THE INVENTION

[0001] The present invention relates to devices for managing voltage levels to computer electronic devices and, in turn, maintains a desired level of performance.

BACKGROUND OF THE INVENTION

[0002] Voltage regulators are widely used in hardware designs to maintain a steady supply of voltage via a power distribution bus to the various components of an exemplary computer system.

[0003] The use of a voltage regulator in the prior art to provide a consistent voltage level at the microprocessor of a computer system is well known. Figure 1 shows a typical voltage regulator as it is used in the prior art in a typical computer system; the voltage regulator(s) provide a consistent voltage level to one or more various components of a computer system via a power line or a power distribution bus. In many systems today, a computer's microprocessor may have its own, dedicated power regulator.

[0004] However, in certain situations the voltage regulator fails to maintain the necessary voltage level resulting in a droop in the power level at the microprocessor, which leads to a delay or failure in performance. One such situation occurs when the microprocessor is at a dormant (e.g. half asleep) state, at which time the microprocessor is called upon to perform a function, which requires a certain amount of power. In some instances the change in current consumption can be as much as 50 amps. This abrupt demand for power from the microprocessor results in a voltage droop in the output of the

1063766.01B03

voltage regulator, which under normal circumstances remains constant. This voltage shortage results in a drop in power level at the microprocessor, which may cause the microprocessor to hang/crash (e.g. the voltage droop reduces the microprocessor's maximum operating frequency) for a period of time until voltage levels are eventually restored. The maximum operating frequency ("Fmax") of a microprocessor will drop as its supply voltage drops. If the microprocessor has an operating frequency of "Foper" and the supply voltage drops such that $F_{oper} > F_{max}$, then the system will hang/crash. Figure 2 illustrates how the voltage droop occurs in the output of the voltage regulator when the microprocessor changes states (e.g. when a large clock domain in the microprocessor is turned on). Two mechanisms can cause short term voltage droop. One is due to inductance, and the other is due to resistance. Inductance voltage droop is described by the equation $V = L \frac{di}{dt}$ where the voltage drop across an inductance L is given by the derivative of the current through the inductance with respect to time. Resistive voltage droop is described by the equation $V = IR$, where I is the current and R is resistance; as I increases, the voltage droop increases.

[0005] Figure 3 shows an example in the prior art of dealing with this problem of power supply droop. One or more levels of capacitors typically support the voltage on the microprocessor die: (a) capacitors on the die; (b) capacitors on the package for the die (around the perimeter of the die); and (c) capacitors on the printed circuit board immediately opposite the microprocessor package. When the voltage droop occurs as in the situation described above, the capacitors discharge, sourcing current, and power is passed to the microprocessor to compensate for the shortage in power due to increased power loading of the microprocessor. However, this technique has significant

disadvantages; the capacitance at (a) has limited size as it is on the die and is very expensive, and capacitance at (b) and (c) is limited by the path impedance from the Core Vdd to the capacitors to Vss and back to the Vss of the Core. For an advanced microprocessor the path inductance needs to be less than 20 pH, and the path resistance needs to be less than 100 micro ohms. This gives a droop level of about an appreciable amount of millivolts due to practical limits. Decreasing the droop magnitude requires exponentially less power supply path inductance. This can be attempted with exotic capacitors and more expensive package/board methods, but they cannot eliminate the effect. The invention eliminates the effect at a cheaper cost point.

[0006] An example of an embodiment of the present invention is to resolve this problem of droop in the power supply, which can lead to a decrease in microprocessor maximum operating frequency.

SUMMARY OF THE INVENTION

[0007] Methods and apparatuses for maintaining power in a data processing system are disclosed. In order to maintain a certain level of power in a data processing system, one apparatus of the present invention, in an exemplary embodiment, detects power level change in the system by detecting a change in current or voltage. In this embodiment, once the power level change is detected, a signal is relayed to a component (e.g. a voltage controller) to boost voltage to a processor (e.g. a microprocessor) through a power distribution bus and as such compensate for a potential voltage droop. An embodiment of one method comprises using a current detector to ascertain changes in power by detecting change in current drawn by a microprocessor, and as such, determine whether to increase voltage to a certain level for a predetermined amount of time to compensate for any voltage droop and, in turn, maintain power in the data processing system.

[0008] In an alternative embodiment, change in power is detected by programming a processor (e.g. a microprocessor) to anticipate changes in compute load and as such determine the power level needed to maintain the processor's frequency by instructing, in one embodiment, a voltage controller to increase voltage to compensate for any voltage droop that may occur at the processor (e.g. microprocessor).

[0009] In another exemplary embodiment, power loading can be monitored to determine the computational load, and this information can be used to select operating states. The states would be characterized by voltage and operating frequency. For example, a given operating voltage would have a corresponding F_{max} (e.g. a low voltage "V_{low}" would correspond to a maximum operating frequency of "F_{maxlow}," and a higher operating voltage "V_{high}" would have a higher maximum operating frequency of

“Fmaxhigh”). When compute load is low (for example, during a time when the computer system is in an idle state), the microprocessor and/or other hardware is operated at a lower frequency than the “Fmaxlow” frequency (operating frequency < “Fmaxlow”) and the voltage applied to the microprocessor and/or other hardware is set to “Vlow.” When the computational load increases (e.g. a large, previously idle portion of the microprocessor starts to consume greater power than before), the system would move to a higher voltage state so that the system can then support (after the higher voltage state is established) a higher operating frequency. The operating frequency is changed to the higher operating frequency after the higher voltage state is established. Once the system reduces current consumption (e.g. a return to an idle state for the portion of the microprocessor), this reduction can be detected and the system’s operating frequency can be reduced and then the supply voltage level is reduced.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] Figure 1 shows a typical prior art voltage regulator as it relates to a computer system.

[0011] Figure 2 shows how the prior art (e.g. the voltage regulator) is unsuccessful in preventing the occurrence of voltage droops.

[0012] Figure 3 shows another typical prior art approach, in which, capacitors are coupled with the microprocessor's Die in order to deal with voltage droop.

[0013] Figure 4 shows an example of an embodiment according to the present invention.

[0014] Figure 5 shows a graph representing how an example of an embodiment according to the present invention increases the voltage to a certain limit in order to compensate for a voltage droop.

[0015] Figure 6 shows how the increase in voltage represented in Figure 5 is maintained for a limited amount of time.

[0016] Figure 7 shows a flow chart showing an exemplary method according to the present invention.

[0017] Figure 8 shows an example of an embodiment according to the present invention, which represents how the present invention is a function of various variables.

[0018] Figure 9 shows a more detailed example of an embodiment according to the present function.

[0019] Figure 10A shows an example of how one embodiment of the invention takes advantage of the relationship between current and compute load in order to manage the voltage level during power (watts) load changes.

[0020] Figure 10B shows a graph representing how the current relates to the voltage needed to maintain a certain power level.

[0021] Figure 11 shows an example of an embodiment of the present invention and how it relates to an exemplary computer system.

[0022] Figure 12 shows another example of a method of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

[0023] In the following description for purposes of explanation, specific systems, interconnections and signals are set forth in order to provide a thorough understanding of the present invention.

[0024] Figure 4 shows one embodiment of the present invention. The voltage controller 400 receives a power input from a power supply source, such as, for example, AC to DC rectifier circuit, and provides as an output a controlled voltage. This output is provided to power distribution planes (in the printed circuit board) which distribute the power to components in the system such as a processor (a microprocessor). These planes are referred to herein as a power distribution bus. It will be appreciated from the discussion herein that the voltage controller 400 does not function in the same manner as a conventional voltage regulator which merely attempts to maintain a constant voltage level regardless of power consumption. A voltage controller 400 is coupled to receive a signal to boost voltage 403 from a detector of power change 401. The voltage controller 400 is coupled with a power distribution bus 402 to provide power to, for example, a microprocessor in the system which is coupled to this bus. The detector of power change 401 detects a change in one of current or voltage and is coupled to send a signal to boost voltage 403 to the voltage controller 400, which is coupled to relay a voltage boost through a power distribution bus 402 (which is typically a plurality of conductive planes in a printed circuit board in one case). The detector of power change 401 detects a change in current drawn by a processor 100 (described below under Figure 11) and is typically coupled to the processor. The current detector may, in one embodiment,

measure current directly or may, in an alternative embodiment, sense magnetic flux or magnetic flux changes.

[0025] Figure 5 shows a representation of an adjustable reference voltage, which voltage controller 400 provides, in one embodiment, in response to a signal which represents a current level. The curve of Figure 5 shows a relationship between an output voltage level from voltage controller 400 and a current level, which is detected by the detector of power (e.g. current) change 401 and is represented in a signal 403. A voltage controller 400 may produce the behavior of the curve of Figure 5, using a look up table (LUT) which maps, through the curve, an input current (I), along the horizontal "X" axis of Figure 5, to an output voltage, along the vertical "Y" axis of Figure 5, to be outputted by the voltage controller 400 in response to the input current. Of course, the values in the LUT may be obtained directly from the curve's abscissa and ordinate values in Figure 5. An example of a method for generating the values of the curve is described herein (e.g. see Figures 10A and 10B). Thus, in one exemplary embodiment, a detector of power change, such as detector 401, senses a change in power load (e.g. an increase in current consumed by a microprocessor) and in turn causes a signal, such as a boost signal 403, to be sent to a voltage controller, such as controller 400 which in turn increases, at least momentarily, the output voltage. The increase in output voltage may be constrained to a maximum amount (e.g. as determined by the safe operating limits for an input voltage to an integrated circuit), such as V_{max} as shown in Figure 5, and the actual increased output voltage may, in one embodiment, be determined by the curve of Figure 5, through a LUT as described above. Thus, in the case of a LUT implementation, signal 403 provides a value representing an I value on the X axis of Figure 5 which maps, through the curve's

values stored in the LUT, to an output voltage value on the "Y" axis. Alternatively, the actual increased output voltage may be a predetermined stepped increase (e.g. see Figure 10A) or merely a single, predetermined increase (regardless of current consumption) which does not require any mapping. The increased output voltage is normally for a momentary period of time which usually coincides with the start up of a large portion of circuitry in an integrated circuit (IC).

[0026] Figure 6 shows a representation, in one embodiment, of how the voltage controller 400, when it receives a boost voltage control signal 403, increases voltage, through the power distribution bus 402, at for example a microprocessor (as described below under Figure 11) for a limited time (represented as $t_1 - t_2$). The curve of Figure 6 illustrates the output voltage from an embodiment of a voltage controller of the invention, such as voltage controller 400. As shown in Figure 6, at times before t_1 and at times after t_2 , the output voltage at a system's microprocessor is at a constant value and thus the voltage controller is acting in a manner which resembles a voltage regulator. At time t_1 a change in power consumption (e.g. an increase in current drawn by the microprocessor) is detected, causing the voltage controller to increase its output voltage, which is supplied to the power distribution bus, from the period of time between t_1 and t_2 . The actual amount of the output voltage increase (between t_1 and t_2) may be a fixed, predetermined amount (which is the same regardless of the change in current drawn by the microprocessor or other circuitry) or it may depend upon the amount of current or power being consumed (e.g. it may depend in a manner reflected by the curve of Figure 5 such that the output voltage increase may be X when $I=A$ and the output voltage increase may

be Y when I=B, where $X>Y$ when $A>B$ (assuming that the output voltage is constrained within V_{min} and V_{max} as shown in Figure 5)).

[0027] Figure 11 shows a typical computer system, which may incorporate the apparatus of the present invention. The computer may include a processor 100, which is typically a microprocessor. The microprocessor is coupled to various computer components through a BUS (or BUSES) 104. Similarly, input and output devices 106, such as keyboards, hard disks, and other well-known computer devices are coupled through one or more input/output controllers 105 which are coupled through a BUS (or BUSES) 104 to other computer system components such as the Memory 101 and Storage devices 102. The Memory and Memory controller 101 as well as Storage Devices (e.g. Hard Drive) 102 are coupled to the computer system components through a BUS (or BUSES) 104. Other components of a typical computer system are also represented in Figure 11 such as Display Devices 107, which are coupled through one or more Display controllers 108 that are coupled to other computer system components through a BUS (or BUSES) 104.

[0028] Figure 11 also shows a particular embodiment of the current invention as it is incorporated in a typical computer system. The apparatus includes a Current Detector 103 and a Voltage Controller 110. The Current Detector is coupled to detect current change in the Microprocessor 100, which is indicative of power change in the Microprocessor. The change in current is also indicative of the compute load level of the microprocessor as represented in Figure 10A, and as the compute load increases, the current drawn increases resulting in a short term voltage droop. The voltage droop may occur due to the sudden demand for power that is the result, for example, of the change in

the microprocessor state from half asleep to full power (e.g. turning on a large clock domain), as described in Figure 2. An uncompensated power supply droop limits the maximum operating frequency (F_{max}) of the microprocessor 100. In many systems, the microprocessor has its own dedicated voltage controller which may provide several controlled voltages to the microprocessor (e.g. 1.2v, 3.3v, etc.).

[0029] The Current Detector 103 is coupled to send a signal to the Voltage Controller 110. Once the Current Detector 103 detects change in the current at the microprocessor 100, it signals the Voltage Controller 110 to boost the voltage for a limited amount of time as represented in Figure 6. The Voltage Controller 110 relays the voltage boost to the microprocessor through the power distribution bus (e.g. the microprocessor's voltage planes). This boost in voltage serves to compensate for any voltage droop that may occur due to the sudden demand for power at the microprocessor and as such enable the microprocessor to handle its compute load without experiencing a lower maximum operating frequency.

[0030] Figure 9 shows a more detailed representation of a voltage controller of the invention such as the Voltage Controller 110. The Current Detector in Figure 9 902 is coupled to send a signal to an Adjustable Reference Voltage device 903. The Adjustable Reference Voltage 903 functions as a translator, which translates current levels to reference voltage levels. Once the Current Detector 902 detects the current level at the microprocessor it relays that information to the Adjustable Reference Voltage 903. In turn, the Adjustable Reference Voltage 903 determines the appropriate voltage level that must be provided by the Voltage Controller 900 according to the current level detected and as such the Voltage Controller boosts the voltage to the appropriate level through the

403766-0189
[0030] differential amplifier of Figure 9. The current to voltage relationship, of one embodiment, is exemplified in Figure 5. It will be appreciated that the variable resistor, shown coupled to the output of the differential amplifier, may be implemented as multiple variable controlled resistors in parallel. An alternative of the controller of Figure 9 may use a switched capacitor to switch in an additional voltage source for a period of time. In this alternative, V_{in} may be an auxiliary voltage on an auxiliary capacitor. When it is desired to add a voltage boost (V_{boost}) to mitigate voltage droop effects, the capacitor is switched into the output of the controller, and the charge on the capacitor would quickly move the V_{out} node, raising the voltage. The amount of the boost would be set by the ratio of the capacitance on V_{out} , the capacitor and the voltage difference prior to coupling/switching the capacitor onto V_{out} .

[0031] An exemplary method of the present invention will now be described by referring to Figure 7. The method begins in 711. In 711, the current detector detects current change to determine whether power demand is increasing or is about to increase. In 712 if no current change is detected the process restarts. However, in 713, if current change is detected (e.g. the current increases rapidly beyond a threshold amount over a given period of time) then voltage through the voltage controller is increased for a limited time and up to a predetermined value. In 714, the voltage level is brought back to its normal level. An example of the behavior of the voltage level over time as a result of operation 713 is shown in Figure 6.

[0032] Figure 7 also shows an alternative method of the present invention, which utilizes programming instructions (e.g. snoop instruction) to program the microprocessor to anticipate power supply droop. In one embodiment of this method, the microprocessor

is programmed to anticipate changes in current by examining instructions that are queued to be executed at the microprocessor and thus anticipate whether power demand is going to increase (e.g. if a co-processor which is asleep is about to be awakened because queued instructions will require the use of the co-processor). In 712 if no current change is detected the process restarts. However, in 713, if current change is detected then voltage through the voltage controller is increased for a limited time and up to a predetermined value. In 714, the voltage level is brought back to its normal level.

[0033] In another exemplary embodiment, power loading can be monitored to determine the computational load, and this information can be used to select operating states. The states would be characterized by voltage and operating frequency. For example, a given operating voltage would have a corresponding F_{max} (e.g. a low voltage "Vlow" would correspond to a maximum operating frequency of " F_{maxlow} ," and a higher operating voltage "Vhigh" would have a higher maximum operating frequency of " $F_{maxhigh}$ "). When compute load is low (for example, during a time when the computer system is in an idle state), the microprocessor and/or other hardware is operated at a lower frequency than the " F_{maxlow} " frequency (operating frequency < " F_{maxlow} ") and the voltage applied to the microprocessor and/or other hardware is set to "Vlow." When the computational load increases (e.g. a large, previously idle portion of the microprocessor starts to consume greater power than before), the system would move to a higher voltage state so that the system can then support (after the higher voltage state is established) a higher operating frequency. The operating frequency is changed to the higher operating frequency after the higher voltage state is established. Once the system reduces current consumption (e.g. a return to an idle state for the portion of the

microprocessor), this reduction can be detected and the system's operating frequency can be reduced and then the supply voltage level is reduced.

[0034] The change from a low power supply voltage to a higher supply voltage may include the additional boost voltage (for a period of time) to mitigate the voltage droop effect. An example of this embodiment, with changes to states in the system, will now be provided while referring to Figure 12. While operating a processing system (e.g. a microprocessor) at a first frequency, a determination is made in operation 1201 of Figure 12, to change to a faster frequency (a "second" frequency). Prior to changing the operating frequency, the power supply voltage is raised in operation 1203 (from V to V'). An optional operation (1205) may then follow in which V' is increased to $V' + V_{\text{boost}}$ for a period of time; this is similar to operation 713 of Figure 7. This additional voltage boost is designed to mitigate the effects of voltage droop. Following operation 1205, the operating frequency is increased to a second frequency in operation 1207. Then, in operation 1209, the power supply voltage is reduced back to V' and the system continues to operate at the second operating frequency. In operation 1211, a determination is made to reduce the operating frequency (e.g. the system automatically decides to enter a reduced power consumption state or the user selects such a state). As a result, in operation 1213, the operating frequency of the system is lowered to the first frequency (or to another frequency which is lower than the second frequency). After operation 1213, the power supply voltage is lowered from V' to V (or to a voltage appropriate to the corresponding new operating frequency).

[0035] Figure 8 shows an embodiment of the current invention represented by a power controller 800, which functions in response to inputs such as current of the system

801, voltage of the system 802, frequency of the system 803 and power supply input 804.

The power controller 800 determines a power level (e.g. of a microprocessor) as a function of the operating frequency of the system 803 (e.g. a data processing system), the voltage of the system 802, and the current of the system 801. The power controller 800, after factoring in the inputs described above, determines the output voltage 805 which is supplied through a power distribution bus (e.g. power distribution planes) in order to maintain a power level (e.g. at the microprocessor) and, in turn, prevent the occurrence of (any) voltage droop or mitigate the effect of a voltage droop.

[0036] Figure 10A shows an exemplary representation of how a system's compute load relates to a system current (I). As compute load increases, for example, a large portion of a microprocessor exits a low power mode and enters a higher power mode, so does the current level, which, in the current invention, is relayed by a current detector to a power or voltage controller of the invention as described above. Figure 10A may be created by empirically testing a computer system at various levels of compute load (e.g. microprocessor completely asleep (e.g. drawing reduced power) or microprocessor partially asleep (e.g. when a numeric co-processor portion of the microprocessor is asleep but the rest of the microprocessor is active, drawing nearly full power) or microprocessor is completely active) and determining, for each given compute load, the corresponding current drawn by the system or a component of the system such as the system's microprocessor(s). Thus, as shown in Figure 10A, a system may be run at three different compute loads (CL1, CL2, and CL3, representing increasing levels of computing resources being used in the system) and the corresponding current levels (I1, I2 and I3) may be measured at these different compute loads. This information may then be used to

derive a curve such as Figure 10B or a curve such as Figure 5 which show the desired output voltage level from a voltage controller of the invention such as voltage controller 400 of Figure 4. For example, after determining the three current levels, I1, I2, and I3 in Figure 10A, the required increase in output voltage from a voltage controller of the invention may be determined. Depending on the system, an increase of this output voltage from about 10 to 200 millivolts, for a short duration (e.g. about several microseconds) may be sufficient to reduce the effects of a voltage droop. The change in voltage from the voltage controller will typically be constrained by Vmin and Vmax (for a power supply voltage) which are set by the IC's physical characteristics, but there is typically a range within which a voltage controller of the invention may operate. The desired increases in output voltage from a voltage controller of the invention for corresponding current levels may be determined empirically for a system by varying the output voltage from such a voltage controller and measuring system performance (e.g. MIPS or operating frequency (MHz) or other parameters or various combinations of parameters). Figure 10B shows an example of the result of such a determination which, as does the curve of Figure 5, maps given current consumption levels to desired output voltage levels from a voltage controller of the invention.

[0037] In the foregoing specification, the invention has been described with reference to specific exemplary embodiments thereof. It will be evident that various modifications may be made thereto without departing from the broader spirit and scope of the invention as set forth in the following claims. The specification and drawings are, accordingly, to be regarded in an illustrative sense rather than a restrictive sense and are not to be construed in a manner which limits the following claims.